arXiv:1410.5088v1 [cs.LO] 19 Oct 2014

Propositional Reasoning about Safety and Termination of Heap-Manipulating Programs

Cristina David, Daniel Kroening, and Matt Lewis

University of Oxford

Abstract. This paper shows that it is possible to reason about the safety and termination of programs handling potentially cyclic, singly-linked lists using propositional reasoning even when the safety invariants and termination arguments depend on constraints over the lengths of lists. For this purpose, we propose the theory SLH of singly-linked lists with length, which is able to capture non-trivial interactions between shape and arithmetic. When using the theory of bit-vector arithmetic as a background, SLH is efficiently decidable via a reduction to SAT. We show the utility of SLH for software verification by using it to express safety invariants and termination arguments for programs manipulating potentially cyclic, singly-linked lists with unrestricted, unspecified sharing. We also provide an implementation of the decision procedure and use it to check safety and termination proofs for several heap-manipulating programs.

Keywords: Heap, SAT, safety, termination.

1 Introduction

Proving safety of heap-manipulating programs is a notoriously difficult task. One of the main culprits is the complexity of the verification conditions generated for such programs. The constraints comprising these verification conditions can be arithmetic (e.g. the value stored at location pointed by x is equal to 3), structural (e.g. x points to an acyclic singly-linked list), or a combination of the first two when certain structural properties of a data structure are captured as numeric values (e.g. the length of the list pointed by x is 3). Solving these combined constraints requires non-trivial interaction between shape and arithmetic.

For illustration, consider the program in Figure 1b, which iterates simultaneously over the lists x and y. The program is safe, i.e., there is no null pointer dereferencing and the assertion after the loop holds. While the absence of null pointer dereferences is trivial to observe and prove, the fact that the assertion after the loop holds relies on the fact that at the beginning of the program and after each loop iteration the lengths of the lists z and t are equal. Thus, the specification language must be capable of expressing the fact that both z and t reach null in the same number of steps. Note that the interaction between shape and arithmetic constraints is intricate, and cannot be solved by a mere theory combination.

The problem is even more pronounced when proving termination of heap-manipulating programs. The reason is that, even more frequently than in the case of safety checking, termination arguments depend on the size of the heap data structures. For example, a loop iterating over the nodes of such a data structure terminates after all the reachable nodes have been explored. Thus, the termination argument is directly linked to the number of nodes in the data structure. This situation is illustrated again by the loop in Figure 1b.

There are few logics capable of expressing this type of interdependent shape and arithmetic constraint. One of the reasons is that, given the complexity of the constraints, such logics can easily become undecidable (even the simplest use of transitive closure leads to undecidability [1]), or at best inefficient.

The tricky part is identifying a logic expressive enough to capture the corresponding constraints while still being efficiently decidable. One work that inspired us in this endeavor is the recent approach by Itzhaky et al. on reasoning about reachability between dynamically allocated memory locations in linked lists using effectively-propositional (EPR) reasoning [2]. This result is appealing as it can harness advances in SAT solvers. The only downside is that the logic presented in [2] is better suited for safety than termination checking, and is best for situations where safety does not depend on the interaction between shape and arithmetic. Thus, our goal is to define a logic that can be used in such scenarios while still being reducible to SAT.

This paper shows that it is possible to reason about the safety and termination of programs handling potentially cyclic, singly-linked lists using propositional reasoning. For this purpose, we present the logic SLH which can express interdependent shape and arithmetic constraints. We empirically prove its utility for the verification of heap-manipulating programs by using it to express safety invariants and termination arguments for intricate programs with potentially cyclic, singly-linked lists with unrestricted, unspecified sharing.

SLH is parametrised by the background arithmetic theory used to express the length of lists (and implicitly every numeric variable). The decision procedure reduces validity of a formula in SLH to satisfiability of a formula in the background theory. Thus, SLH is decidable if the background theory is decidable.

As we are interested in a reduction to SAT, we instantiate SLH with the theory of bit-vector arithmetic, resulting in SLH[$\mathcal{T}_{\mathcal{BV}}$]. This allows us to handle non-linear operations on lists length (e.g. the example in Figure 1c), while still retaining decidability. However, SLH can be combined with other background theories, e.g., Presburger arithmetic.

We provide an implementation of our decision procedure for SLH[$\mathcal{T}_{\mathcal{BV}}$] and test its efficiency by verifying a suite of programs against safety and termination specifications expressed in SLH. Whenever the verification fails, our decision procedure produces a counterexample.

Contributions:
* We propose the theory SLH of singly-linked lists with length. SLH allows *unrestricted* sharing and cycles.
* We define the strongest post-condition for formulae in SLH.
* We show the utility of SLH for software verification by using it to express safety invariants and termination arguments for programs with potentially cyclic singly-linked lists.
* We present the instantiation SLH[$\mathcal{T}_{\mathcal{BV}}$] of SLH with the theory of bit-vector arithmetic. SLH[$\mathcal{T}_{\mathcal{BV}}$] can express non-linear operations on the lengths of lists, while still retaining decidability.
* We provide a reduction from satisfiability of SLH[$\mathcal{T}_{\mathcal{BV}}$] to propositional SAT.
* We provide an implementation of the decision procedure for SLH[$\mathcal{T}_{\mathcal{BV}}$] and test it by checking safety and termination for several heap-manipulating programs (against provided safety invariants and termination arguments).

2 Motivation

Consider the examples in Figure 1. They all capture situations where the safety (i.e. absence of null pointer dereferencing and no assertion failure) and termination of the program depend on interdependent shape and arithmetic constraints. In this section we only give an intuitive description of these examples, and we revisit and formally specify them in Section 7. We assume the existence of the following two functions: (1) $length(x)$ returns the number of nodes on the path from x to NULL if the list pointed by x is acyclic, and MAXINT otherwise; (2) $circular(x)$ returns true iff the list pointed by x is circular (i.e., x is part of a cycle).

In Figure 1a, we iterate over the potentially cyclic singly-linked list pointed by x a number of times equal with the result of $length(x)$. The program is safe (i.e. y is not NULL at loop entry) and terminating. A safety invariant for the loop needs to capture the length of the path from y to NULL.

The loop in Figure 1b iterates over the lists pointed by x and y, respectively, until one of them becomes NULL. In order to check whether the assertion after the loop holds, the safety invariant must relate the length of the list pointed by x to the length of the list pointed by y. Similarly, a termination argument needs to consider the length of the two lists.

The example in Figure 1c illustrates how non-linear arithmetic can be encoded via singly-linked lists. Thus, the loop in $divides(x, y)$ iterates over the list pointed by x a number of nodes equal to the quotient of the integer division $length(x)/length(y)$ such that, after the loop, the list pointed by z has a length equal with the remainder of the division.

The function in Figure 1d returns true iff the list passed in as a parameter is circular. The functional correctness of this function is captured by the assertion

```
List x, y = x;
int n = length(x), i = 0;

while (i < n) {
  y = y→next;
  i = i+1;
}
```

(a)

```
List x, y, z = x, t = y;

assume(length(x)==length(y));

while (z!=NULL && t!=NULL) {
  z = z→next;
  t = t→next;
}

assert (z==NULL && t==NULL);
```

(b)

```
int divides(List x, List y) {
  List z = y;
  List w = x;

  assume(length(x)!=MAXINT &&
         length(y)!=MAXINT &&
         y != NULL);

  while (w != NULL) {
    if (z == NULL) z = y;
    z = z→next;
    w = w→next;
  }

  assert(z == NULL ⇔
         length(x)%length(y)==0);
  return z == NULL;
}
```

(c)

```
int isCircular(List l) {
  List p = q = l;

  do {
    if (p != NULL) p = p→next;
    if (q != NULL) q = q→next;
    if (q != NULL) q = q→next;
  }
  while (p!=NULL && q!=NULL
         && p!=q);

  assert(p == q ⇔ circular(l));
  return p == q;
}
```

(d)

Fig. 1: Motivational examples.

after the loop checking that pointers p and q end up being equal iff the list l is circular.

3 Theory of Singly Linked Lists with Length

In this section we introduce the theory SLH for reasoning about potentially cyclic singly linked lists.

3.1 Informal Description of SLH

We imagine that there is a set of pointer variables $x, y, \ldots$ which point to heap cells. The cells in the heap are arranged into singly linked lists, i.e. each cell has a "next" pointer which points somewhere in the heap. The lists can be cyclic and two lists can share a tail, so for example the following heap is allowed in our logic:

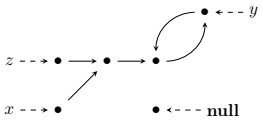

Our logic contains functions for examining the state of the heap, along with the four standard operations for mutating linked lists: *new*, *assign*, *lookup* and *update*. We capture the side-effects of these mutation operators by explicitly naming the current heap – we introduce heap variables h, h' etc. which denote the heap in which each function is to be interpreted. The mutation operators then become pure functions mapping heaps to heaps. The heap functions of the logic are illustrated by example in Figure 3 and have the following meanings:

$alias(h, x, y)$: do x and y point to the same cell in heap h?

$isPath(h, x, y)$: is there a path from x to y in h?

$pathLength(h, x, y)$: the length of the shortest path from x to y in h.

$isNull(h, x)$: is x **null** in h?

$circular(h, x)$: is x part of a cycle, i.e. is there some non-empty path from x back to x in h?

$h' = new(h, x)$: obtain h' from h by allocating a new heap cell and reassigning x so that it points to this cell. The newly allocated cell is not reachable from any other cell and its successor is **null**. This models the program statement $x = new()$. For simplicity, we opt for this allocation policy, but we are not restricted to it.

$h' = assign(h, x, y)$: obtain h' from h by assigning x so that it points to the same cell as y. Models the statement $x = y$.

$h' = lookup(h, x, y)$: obtain h' from h by assigning x to point to y's successor. Models the statement $x = y{\rightarrow}next$.

$h' = update(h, x, y)$: obtain h' from h by updating x's successor to point to y. Models $x{\rightarrow}next = y$.

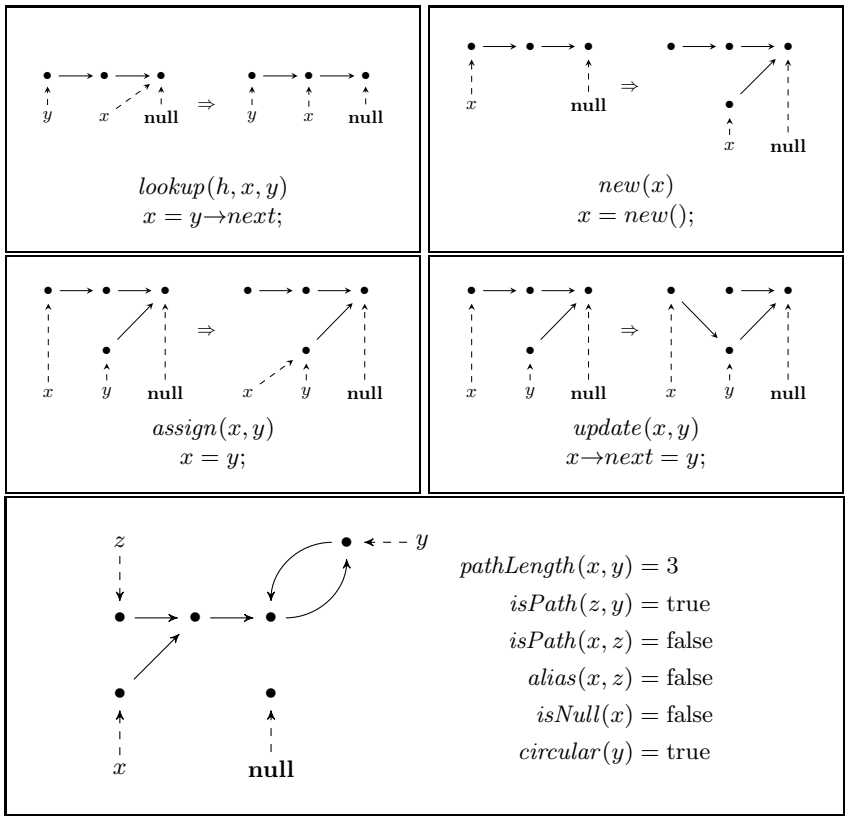

Fig. 3: SLH by example

3.2 Syntax of SLH

The theory of singly-linked lists with length, SLH, uses a background arithmetic theory $\mathcal{T}_{\mathcal{B}}$ for the length of lists (implicitly any numeric variable). Thus, SLH has the following signature:

$$\begin{aligned}\Sigma_{SLH} = \Sigma_B \cup \{&\mathit{alias}(\cdot,\cdot,\cdot), \mathit{isPath}(\cdot,\cdot,\cdot), \mathit{pathLength}(\cdot,\cdot,\cdot), \mathit{isNull}(\cdot,\cdot),\\ &\mathit{circular}(\cdot,\cdot), \cdot{=}\mathit{new}(\cdot,\cdot), \cdot{=}\mathit{assign}(\cdot,\cdot,\cdot),\\ &\cdot{=}\mathit{lookup}(\cdot,\cdot,\cdot), \cdot{=}\mathit{update}(\cdot,\cdot,\cdot)\}.\end{aligned}$$

where the nine new symbols correspond to the heap-specific functions described in the previous section.

Sorts. Heap variables (e.g. h in $\mathit{alias}(h,x,y)$) have sort $\mathcal{S}_{\mathcal{H}}$, pointer variables have sort $\mathcal{S}_{Addr}$ (e.g. x and y in $\mathit{alias}(h,x,y)$), numeric variables have sort $\mathcal{S}_{\mathcal{B}}$ (e.g. n in $n = \mathit{pathLength}(h,x,y)$).

Literal and formula. A literal in SLH is either a heap function or a $\mathcal{T}_{\mathcal{B}}$-literal. A formula in SLH is a boolean combination of SLH-literals.

3.3 Semantics of SLH

We give the semantics of SLH by defining the models in which an SLH formula holds. An interpretation Γ is a function mapping free variables to elements of the appropriate sort. If an SLH formula ϕ holds in some interpretation Γ, we say that Γ *models* ϕ and write $\Gamma \models \phi$.

Interpretations may be constructed using the following substitution rule:

$$\Gamma[h \mapsto H](x) = \begin{cases} H & \text{if } x = h \\ \Gamma(x) & \text{otherwise} \end{cases}$$

Pointer variables are considered to be a set of constant symbols and are thus given a fixed interpretation. The only thing that matters is that their interpretation is pairwise different. We assume that the pointer variables include a special name **null**. The set of pointer variables is denoted by the symbol P.

We will consider the semantics of propositional logic to be standard and the semantics of $\mathcal{T}_{\mathcal{B}}$ given, and thus just define the semantics of heap functions. To do this, we will first define the class of objects that will be used to interpret heap variables.

Definition 1 (Heap). *A heap over pointer variables P is a pair $H = \langle L, G\rangle$. G is a finite graph with vertices $V(G)$ edges $E(G)$. $L : P \to V(G)$ is a labelling function mapping each pointer variable to a vertex of G. We define the cardinality of a heap to be the cardinality of its underlying graph's vertices: $\|H\| = \|V(G)\|$.*

Definition 2 (Singly Linked Heaps). *A heap $H = \langle L, G\rangle$ is a singly linked heap iff each vertex has outdegree 1, except for a single sink vertex that has outdegree 0 and is labelled by* **null***:*

$$\forall v \in V(G).(\text{outdegree}(v) = 1 \wedge L(\mathbf{null}) \neq v) \vee (\text{outdegree}(v) = 0 \wedge L(\mathbf{null}) = v)$$

Having defined our domain of discourse, we are now in a position to define the semantics of the various heap functions introduced in Section 3.1. We begin with the functions examining the state of the heap and will use a standard structural recursion to give the semantics of the functions with respect to an implicit interpretation Γ, so that $[\![h]\!] = \Gamma(h)$. We will use the shorthand $u \xrightarrow{n} v$ to say that if we start at node u, then follow n edges, we arrive at v. We also use $L(H)$ to select the labelling function L from H:

$$u \xrightarrow{n} v \stackrel{\text{def}}{=} \langle u, v\rangle \in E^n$$
$$u \rightarrow^* v \stackrel{\text{def}}{=} \exists n \geq 0. u \xrightarrow{n} v$$
$$u \rightarrow^+ v \stackrel{\text{def}}{=} \exists n > 0. u \xrightarrow{n} v$$

Note that $u \xrightarrow{0} u$. The semantics of the heap functions are then:

$$[\![\mathit{pathLength}(h, x, y)]\!]\Gamma \stackrel{\text{def}}{=} \min\left(\{n \mid L([\![h]\!])(x) \xrightarrow{n} L([\![h]\!])(y)\} \cup \{\infty\}\right)$$
$$[\![\mathit{circular}(h, x)]\!]\Gamma \stackrel{\text{def}}{=} \exists v \in V([\![h]\!]).L([\![h]\!])(x) \rightarrow^+ v \wedge v \rightarrow^+ L([\![h]\!])(x)$$
$$[\![\mathit{alias}(h, x, y)]\!]\Gamma \stackrel{\text{def}}{=} [\![\mathit{pathLength}(h, x, y)]\!] = 0$$
$$[\![\mathit{isPath}(h, x, y)]\!]\Gamma \stackrel{\text{def}}{=} [\![\mathit{pathLength}(h, x, y)]\!] \neq \infty$$
$$[\![\mathit{isNull}(h, x)]\!]\Gamma \stackrel{\text{def}}{=} [\![\mathit{pathLength}(h, x, \mathbf{null})]\!] = 0$$

Note that since the graph underlying H has outdegree 1, *pathLength* and *circular* can be computed in $O(\|H\|)$ time, or equivalently they can be encoded with $O(\|H\|)$ arithmetic constraints.

To define the semantics of the mutation operations, we will consider separately the effect of each mutation on each component of the heap – the labelling function L, the vertex set V and the edge set E. Where a mutation's effect on some heap component is not explicitly stated, the effect is id. For example, *assign* does not modify the vertex set, and so $\mathit{assign}_V = \text{id}$. In the following definitions, we will say that $\text{succ}(v)$ is the unique vertex such that $(v, \text{succ}(v)) \in E(H)$.

$$
\begin{aligned}
[\![new_V(h,x)]\!]\Gamma &\stackrel{\text{def}}{=} V([\![h]\!]) \cup \{q\} \text{where } q \text{ is a fresh vertex} \\
[\![new_E(h,x)]\!]\Gamma &\stackrel{\text{def}}{=} E([\![h]\!]) \cup \{(q, \mathbf{null})\} \\
[\![new_L(h,x)]\!]\Gamma &\stackrel{\text{def}}{=} L([\![h]\!])[x \mapsto q] \\
[\![assign_L(h,x,y)]\!]\Gamma &\stackrel{\text{def}}{=} L([\![h]\!])[x \mapsto L([\![h]\!])(y)] \\
[\![lookup_L(h,x,y)]\!]\Gamma &\stackrel{\text{def}}{=} L([\![h]\!])[x \mapsto \text{succ}(y)] \\
[\![update_E(h,x,y)]\!]\Gamma &\stackrel{\text{def}}{=} (E([\![h]\!]) \setminus \{(L([\![h]\!])(x), \text{succ}(L([\![h]\!])(x)))\}) \cup \\
&\qquad \{(L([\![h]\!])(x), L([\![h]\!])(y))\}
\end{aligned}
$$

4 Deciding Validity of SLH

We will now turn to the question of deciding the validity of an SLH formula, that is for some formula ϕ we wish to determine whether ϕ is a tautology or if there is some Γ such that $\Gamma \models \neg\phi$. To do this, we will show that SLH enjoys a finite model property and that the existence of a fixed-size model can be encoded directly as an arithmetic constraint.

Our high-level strategy for this proof will be to define progressively coarser equivalence relations on SLH heaps that respect the transformers and observation functions. The idea is that all of the heaps in a particular equivalence class will be equivalent in terms of the SLH formulae they satisfy. We will eventually arrive at an equivalence relation (homeomorphism) that is sound in the above sense and which is also guaranteed to have a small heap in each equivalence class.

From here on we will slightly generalise the definition of a singly linked heap and say that the underlying graph is weighted with weight function $w : E(H) \to \mathbb{N}$. When we omit the weight of an edge (as we have in all heaps up until now), it is to be understood that the edge's weight is 1.

4.1 Sound Equivalence Relations

We will say that an equivalence relation $\approx$ is *sound* if the following conditions hold for each pair of pointer variables x, y and transformer τ:

$$
\begin{aligned}
\forall H, H' \cdot H \approx H' \Rightarrow & \mathit{pathLength}(H,x,y) = \mathit{pathLength}(H',x,y) \wedge && (1) \\
& \mathit{circular}(H,x) = \mathit{circular}(H',x) \wedge && (2) \\
& \tau(H) \approx \tau(H') && (3)
\end{aligned}
$$

The first two conditions say that if two heaps are in the same equivalence class, there is no observation that can distinguish them. The third condition says that the equivalence relation is inductive with respect to the transformers. There is therefore no sequence of transformers and observations that can distinguish two heaps in the same equivalence class.

We begin defining two sound equivalence relations:

Definition 3 (Reachable Sub-Heap). *The reachable sub-heap $H|_P$ of a heap H is H with vertices restricted to those reachable from the pointer variables P:*

$$V(H|_P) = \{v \mid \exists p \in P. \langle L(p), v\rangle \in E^*\}$$

Then the relation $\{\langle H, H'\rangle \mid H|_P = H'|_P\}$ is sound.

Definition 4 (Heap Isomorphism). *Two heaps $H = \langle L, G\rangle, H' = \langle G', L'\rangle$ are isomorphic (written $H \simeq H'$) iff there exists a graph isomorphism $f : G|_P \to G'|_P$ that respects the labelling function, i.e., $\forall p \in P. f(L(p)) = L'(p)$.*

Example 1. H and H' are not isomporphic, even though their underlying graphs are.

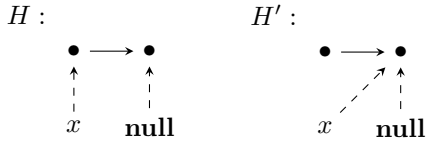

Theorem 1. *Heap isomorphism is a sound equivalence relation.*

4.2 Heap Homeomorphism

The final notion of equivalence we will describe is the weakest. Loosely, we would like to say that two heaps are equivalent if they are "the same shape" and if the shortest distance between pointer variables is the same. To formalise this relationship, we will be using an analog of topological homeomorphism.

Definition 5 (Edge Subdivision). *A graph G' is a subdivison of G iff G' can be obtained by repeatedly subdividing edges in G, i.e., for some edge $(u, v) \in E(G)$ introducing a fresh vertex q and replacing the edge (u, v) with edges $(u, q), (q, v)$ such that $w'(u, q) + w'(q, v) = w(u, v)$. Subdivision for heaps is defined in terms of their underlying graphs.*

We define a function *subdivide*, which subdivides an edge in a heap. As usual, the function is defined componentwise on the heap:

$$\begin{aligned}
\mathit{subdivide}_V(H, u, v, k) &= V \cup \{q\} \\
\mathit{subdivide}_E(H, u, v, k) &= (E \setminus \{(u, v)\}) \cup \{(u, q), (q, v)\} \\
\mathit{subdivide}_W(H, u, v, k) &= W(H)[(u, v) \mapsto \infty, (u, q) \mapsto k, (q, v) \mapsto W(H)(u, v) - k]
\end{aligned}$$

The dual of edge subdivision is edge *smoothing* – if we have two edges $u \xrightarrow{n} q \xrightarrow{m} v$, where q is unlabelled and has no other incoming edges, we can remove q and add the single edge $u \xrightarrow{n+m} v$.

Example 2. H' is a subdivision of H.

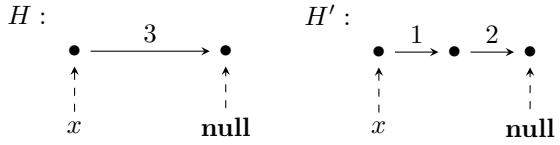

Lemma 1 (Subdividing an Edge Preserves Observations). *If H' is obtained from H by subdividing one edge, then for any x, y we have:*

$$pathLength(H, x, y) = pathLength(H', x, y) \tag{4}$$
$$circular(H, x) = circular(H', x) \tag{5}$$

Definition 6 (Heap Homeomorphism). *Two heaps H, H' are homeomorphic (written $H \sim H'$) iff there there is a heap isomorphism from some subdivision of H to some subdivision of H'.*

Intuitively, homeomorphisms preserve the topology of heaps: if two heaps are homeomorphic then they have the same number of loops and the same number of "joins" (vertices with indegree ≥ 2).

Example 3. H and H' are homeomorphic, since they can each be subdivided to produce S.

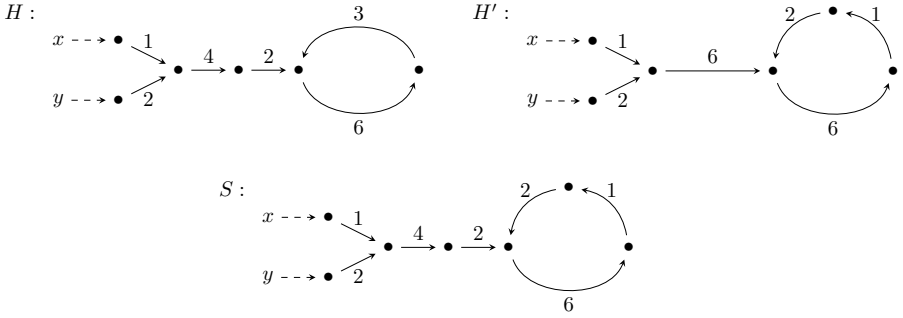

Lemma 2 (Transformers Respect Homeomorphism). *For any heap transformer τ, if $H_1 \sim H_2$ then $\tau(H_1) \sim \tau(H_2)$.*

Proof. It suffices to show that for any transformer τ and single-edge subdivision s, the following diagram commutes:

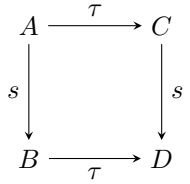

We will check that $\tau \circ s = s \circ \tau$ by considering the components of each arrow separately and using the semantics defined in Section 3.3. The only difficult case is for *lookup*, for which we provide the proof in full. This case is illustrative of the style of reasoning used for the proofs of the other transformers, which can be found in the extended version of this paper.

$\tau = lookup$*:* Now that we have weighted heaps, there are two cases for *lookup*: if the edge leaving $L(y)$ does not have weight 1, we need to first subdivide so that it does; otherwise the transformer is exactly as in the unweighted case, which can be seen easily to commute.

In the second (unweighted) case, all of the components commute due to id. Otherwise, *lookup* is a composition of some subdivision s' and then unweighted lookup: $lookup = lookup_U \circ s'$.

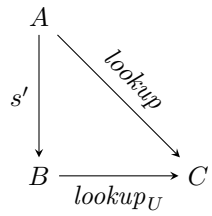

Our commutativity condition is then:

$$(lookup_U \circ s') \circ s = s \circ (lookup_U \circ s')$$

We know that unweighted *lookup* commutes with arbitrary subdivisions, so

$$(lookup_U \circ s') \circ s = s \circ (s' \circ lookup_U)$$
$$lookup_U \circ (s' \circ s) = (s \circ s') \circ lookup_U$$

But the composition of two subdivisions is a subdivision, so we are done.

Theorem 2. *Homeomorphism is a sound equivalence relation.*

Proof. This is a direct consequence of Lemma 1 and Lemma 2.

We would now like to show that for each equivalence class induced by $\sim$, there is a unique minimal element. This is easy to show if we consider the category **SLH** of singly linked heaps, with edge subdivisions as arrows. To recap the definition of a category:

- A *category* is a collection of objects along with a collection of arrows, which are maps from one object to another. We can compose any two arrows $A \to B \to C$ to form one arrow $A \to C$. Composition is associative and each object has an arrow $A \xrightarrow{\text{id}} A$.
- A diagram is a picture consisting of some objects along with some arrows between the objects. We say that the diagram *commutes* if for any two paths between two objects, the arrow generated by composing all the arrows on the first path equals the composition of all the arrows on the second path.

As mentioned, objects of the category **SLH** are singly linked heaps, and there is an arrow from one heap to another if the first can be subdivided into the second. To illustrate, Example 3 is represented in **SLH** by the following diagram:

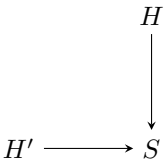

Now for every pair of homeomorphic heaps $H_1 \sim H_2$ we know that there is some X that is a subdivision of both H_1 and H_2. Clearly if we continue subdividing edges, we will eventually arrive at a heap where every edge has weight 1, at which point we will be unable to subdivide any further. Let us call this maximally subdivided heap the *shell*, which we will denote by $\mathrm{Sh}(H_1)$. Then $\mathrm{Sh}(H_1) = \mathrm{Sh}(H_2)$ is the pushout of the previous diagram. Dually, there is some Y that both H_1 and H_2 are subdivisions of, and the previous diagram has a pullback, which we shall call the *kernel*. This is the heap in which all edges have been smoothed. The following diagram commutes, and since a composition of subdivisions and smoothings is a homeomorphism, all of the arrows (and their inverses) in this diagram are homeomorphisms. In fact, the $H_1, H_2, X, Y, \mathrm{Sh}$ and Ke are exactly an equivalance class:

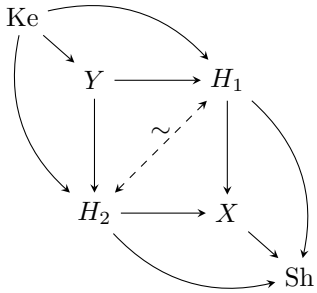

Lemma 3 (Kernels are Small). *For any H, $\|\mathrm{Ke}(H)\| \leq 2 \times \|P\|$.*

Proof. Since $\mathrm{Ke}(H)$ is maximally smoothed, every unlabelled vertex has indegree ≥ 2. We will partition the vertices of H into named and unlabelled vertices:

$$\begin{aligned} N &= \{v \in V(H) \mid \exists p \in P.L(p) = v\} \\ U &= \{u \in V(H) \mid \forall p \in P.L(p) \neq u\} \\ V(H) &= N \cup U \end{aligned}$$

Then let $n = \|N\|$ and $u = \|U\|$. Now, the total indegree of the underlying graph must be equal to the total outdegree, so:

$$\begin{aligned} \sum_{v \in V(H)} \mathrm{out}(v) &= \sum_{v \in V(H)} \mathrm{in}(v) \\ n + u &= \sum_{n \in N} \mathrm{in}(n) + \sum_{u \in U} \mathrm{in}(u) \\ &= \sum_{n \in N} \mathrm{in}(n) + 2u + k \end{aligned}$$

where $k \geq 0$, since $\text{in}(u) \geq 2$ for each u.

$$n = u + \underbrace{\sum_{n \in N} \text{in}(n) + k}_{\geq 0}$$

$$n \geq u$$

So $u \leq n \leq \|P\|$, hence $\|\text{Ke}(H)\| = n + u \leq 2 \times \|P\|$.

Theorem 3 (SLH has Small Model). *For any SLH formula $\forall h.\phi$, if there is a counterexample $\Gamma \models \neg\phi$ then there is $\Gamma' \models \neg\phi$ with every heap-sorted variable in Γ being interpreted by a homeomorphism kernel.*

Proof. This follows from Theorem 2 and Lemma 3.

We can encode the existence of a small model with an arithmetic constraint whose size is linear in the size of the SLH formula, since each of the transformers can be encoded with a constant sized constraint and the observation functions can be encoded with a constraint of size $O(\|H\|) = O(\|P\|)$. An example implementation of the constraints used to encode each atom is given in Section 6. We need one constraint for each of the theory atoms, which leaves us with $O(\|P\| \times \|\phi\|)$ constraints.

Corollary 1 (Decidability of SLH). *If the background theory $\mathcal{T}_\mathcal{B}$ is decidable, then SLH is decidable.*

Proof. The existence of a small model can be encoded with a linear number of arithmetic constraints in $\mathcal{T}_\mathcal{B}$.

5 Using SLH for Verification

Our intention is to use SLH for reasoning about the safety and termination of programs with potentially cyclic singly-linked lists:

- For safety, we annotate loops with safety invariants and generate VCs checking that each loop annotation is genuinely a safety invariant, i.e. (1) it is satisfied by each state reachable on entry to the loop, (2) it is inductive with respect to the program's transition relation, and (3) excludes any states where an assertion violation takes place (the assertions include those ensuring memory safety). The existence of a safety invariant corresponds to the notion of partial correctness: no assertion fails, but the program may never stop running.
- For termination, we provide ranking functions for each loop and generate VCs to check that the loops do terminate, i.e. the ranking function is monotonically decreasing with respect to the loop's body and (2) it is bounded from below. By combining these VCs with those generated for safety, we create a total-correctness specification.

$$
\begin{array}{lll}
\mathit{datat} & := & \mathit{struct}\ C\ \{(\mathit{typ}\ v)^*\} \\
e & := & v \mid v{\rightarrow}\mathit{next} \mid \mathrm{new}(C) \mid \mathbf{null} \\
S & := & v{=}e \mid v_1{\rightarrow}\mathit{next}{=}v_2 \mid S_1; S_2 \mid \mathit{if}\ (B)\ S_1\ \mathit{else}\ S_2 \mid \\
 & & \mathit{while}\ (B)\ S \mid \mathrm{assert}(\phi) \mid \mathrm{assume}(\phi)
\end{array}
$$

Fig. 4: Programming Language

The two additional items we must provide in order to be able to generate these VCs, are a programming language and the strongest post-condition for formulae in SLH with respect to statements in the programming language. We do so next.

5.1 Programming Language

We use the sequential programming language in Fig. 4. It allows heap allocation and mutation, with v denoting a variable and *next* a pointer field. To simplify the presentation, we assume each data structure has only one pointer field, *next*, and allow only one-level field access, denoted by $v{\rightarrow}\mathit{next}$. Chained dereferences of the form $v{\rightarrow}\mathit{next}{\rightarrow}\mathit{next}\ldots$ are handled by introducing auxiliary variables. The statement $\mathit{assert}(\phi)$ checks whether ϕ (expressed in the heap theory described in Section 3) holds for the current program state, whereas $\mathit{assume}(\phi)$ constrains the program state.

For ease of use, when using SLH in the context of safety and termination verification, the SLH functions we expose in the specification language are side-effect free. That is to say, we don't require the explicit heap h to be mentioned in the specifications.

5.2 Strongest Post-Condition

To create a verification condition from a specification, we first decompose the specification into Hoare triples and then compute the strongest post-condition to generate a VC in the SLH theory. Since SLH includes primitive operations for heap manipulation, our strongest post-condition is easy to compute:

$$
\begin{aligned}
\mathsf{SP}(x = y, \phi) &\stackrel{\text{def}}{=} h' = \mathit{fresh}() \wedge \phi[h'/h] \wedge h = \mathit{assign}(h', x, y) \\
\mathsf{SP}(x = y{\rightarrow}\mathit{next}, \phi) &\stackrel{\text{def}}{=} h' = \mathit{fresh}() \wedge \phi[h'/h] \wedge h = \mathit{lookup}(h', x, y) \\
\mathsf{SP}(x = \mathrm{new}(C), \phi) &\stackrel{\text{def}}{=} h' = \mathit{fresh}() \wedge \phi[h'/h] \wedge h = \mathit{new}(h', x, y) \\
\mathsf{SP}(x{\rightarrow}\mathit{next} = y, \phi) &\stackrel{\text{def}}{=} h' = \mathit{fresh}() \wedge \phi[h'/h] \wedge h = \mathit{update}(h', x, y)
\end{aligned}
$$

Where $\mathit{fresh}()$ introduces a fresh heap variable. The remaining cases for SP are standard.

5.3 VC Generation Example

```
x = y;

while (x ≠ null) {
  {isPath(y, x)}
  x = x→next;
}

assert (isPath(y, x));
```

Fig. 5: An annotated program.

Consider the program in Figure 5 which has been annotated with a loop invariant. In order to verify the partial-correctness condition that the assertion cannot fail, we must check the following Hoare triples:

$$\{\top\}\ \mathrm{x} = \mathrm{y}\ \{isPath(y, x)\} \tag{6}$$

$$\{isPath(y, x) \wedge \neg isNull(x)\}\ \mathrm{x} = \mathrm{x} \rightarrow \mathrm{next}\ \{isPath(y, x)\} \tag{7}$$

$$\{isPath(y, x) \wedge isNull(x)\}\ \mathbf{skip}\ \{isPath(y, x)\} \tag{8}$$

Taking strongest post-condition across each of these triples generates the following SLH VCs:

$$\forall h.h' = assign(h, x, y) \Rightarrow isPath(h', y, x) \tag{9}$$

$$\forall h.isPath(h, y, x) \wedge \neg isNull(x) \wedge h' = lookup(h, x, x) \Rightarrow isPath(h', y, x) \tag{10}$$

$$\forall h.isPath(h, y, x) \wedge isNull(x) \Rightarrow isPath(h, y, x) \tag{11}$$

6 Implementation

For our implementation, we instantiate SLH with the theory of bit-vector arithmetic. Thus, according to Corollary 1, the resulting theory SLH[$\mathcal{T}_{\mathcal{BV}}$] is decidable. In this section, we provide details about the implementation of the decision procedure via a reduction to SAT.

To check validity of an SLH[$\mathcal{T}_{\mathcal{BV}}$] formula ϕ we search for a small counterexample H. By Theorem 3, if no such small H exists, there is no counterexample and so ϕ is a tautology. We encode the existence of a small counterexample by constructing a SAT formula.

To generate the SAT formula, we instantiate every occurence of the SLH[$\mathcal{T}_{\mathcal{BV}}$] functions with the functions shown in Figure 6. The structure that the functions operate over is the following, where N is the number of vertices in the structure and P is the number of program variables:

```
typedef int node;
typedef int ptr;

struct heap {
  ptr: node[P];
  succ : (node × int)[N];
  nnodes: int;
}
```

The heap contains N nodes, of which nnodes are allocated. Pointer variables are represented as integers in the range $[0, P-1]$ where by convention **null** $= 0$. Each pointer variable is mapped to an allocated node by the ptr array, with the restriction that **null** maps to node 0. The edges in the graph are encoded in the succ array where h.succ[n] = (m, w) iff the edge (n, m) with weight w is in the graph. For a heap with N nodes, this structure requires $3N + 1$ integers to encode.

The implementations of the SLH[$\mathcal{T}_{\mathcal{BV}}$] functions described in Section 3.1 are given in Figure 6. Note that only `Alloc` and `Lookup` can allocate new nodes. Therefore if we are searching for a counterexample heap with at most $2P$ nodes, and our formula contains k occurences of `Alloc` and `Lookup`, the largest heap that can occur in the counterexample will contain no more than $2P + k$ nodes. We can therefore encode all of the heaps using $6P + 3k + 1$ integers each.

When constructing the SAT formula corresponding to the SLH[$\mathcal{T}_{\mathcal{BV}}$] formula, each of the functions can be encoded (via symbolic execution) as a formula in the background theory $\mathcal{T}_{\mathcal{B}}$ of constant size, except for `PathLength` which contains a loop. This loop iterates $N = 2P + k$ times and so expands to a formula of size $O(P)$. If the SLH[$\mathcal{T}_{\mathcal{BV}}$] formula contains x operations, the final SAT formula in $\mathcal{T}_{\mathcal{BV}}$ is therefore of size $x \times P$. We use CBMC to construct and solve the SAT formula.

One important optimisation when constructing the SAT formula involves a symmetry reduction on the counterexamples. Since our encoding assigns names to each of the vertices in the graph, we can have multiple representations for heaps that are isomorphic. To ensure that the SAT solver only considers a single counterexample from each homeomorphism class, we choose a canonical representative of each class and add a constraint that the counterexample we are looking for must be one of these canonical representatives. We define the canonical form of a heap such that the nodes are ordered topologically and so that the ordering is compatible with the ordering on the progam variables:

$$\forall p, p' \in P.p < p' \Rightarrow \forall n, n'.L(p) \rightarrow^* n \wedge L(p') \rightarrow^* n' \Rightarrow n \leq n'$$
$$\forall n, n'.n \rightarrow n' \Rightarrow n \leq n'$$

Where $n \rightarrow^* n'$ means n' is reachable from n.

```
function NewNode(heap h)
    n ← h.nnodes
    h.nnodes ← h.nnodes + 1
    h.succ[n] ← (null, 1)
    return n

function Subdivide(heap h, node a)
    n ← NewNode(h)
    (b, w) ← h.succ[a]
    h.succ[a] ← (n, 1)
    h.succ[n] ← (b, w - 1)
    return n

function Update(heap h, ptr x, ptr y)
    n ← h.ptr[x]
    m ← h.ptr[y]
    h.succ[n] ← (m, 1)

function Assign(heap h, ptr x, ptr y)
    h.ptr[x] ← h.ptr[y]

function Lookup(heap h, ptr x, ptr y)
    n ← h.ptr[y]
    (n', w) ← h.succ[n]
    if w ≠ 1 then
        n' ← Subdivide(h, n)
    h.ptr[x] ← n'

function Alloc(heap h, ptr x)
    n ← NewNode(h)
    h.ptr[x] ← n

function PathLength(heap h, ptr x, ptr y)
    n ← h.ptr[x]
    m ← h.ptr[y]
    distance ← 0
    for i ← 0 to h.nnodes do
        if n = m then
            return distance
        else
            (n, w) ← h.succ[n]
            distance ← distance + w
    return ∞

function Circular(heap h, ptr x)
    n ← h.ptr[x]
    m ← h.succ[n]
    distance ← 0
    for i ← 0 to h.nnodes do
        if m = n then
            return True
        else
            if n = null then
                return False
        m ← h.succ[m]
    return False
```

Fig. 6: Implementation of the SLH[$\mathcal{T}_{\mathcal{BV}}$] functions

7 Motivation Revisited

In this section, we get back to the motivational examples in Figure 1 and express their safety invariants and termination arguments in SLH. As mentioned in Section 5.1, for ease of use, we don't mention the explicit heap h in the specifications.

In Figure 1a, assuming that the call to the *length* function ensures the state before the loop to be $\mathit{pathLength}(h, x, \mathbf{null}) = n$, then a possible safety invariant is $\mathit{pathLength}(h, y, \mathbf{null}) = n - i$. Note that this invariant covers both the case where the list pointed by x is acyclic and the case where it contains a cycle. In the latter scenario, given that $\infty - i = \infty$, the invariant is equivalent to $\mathit{pathLength}(h, y, \mathbf{null}) = \infty$. A ranking function for this program is $R(i) = -i$.

The program in Figure 1b is safe with a possible safety invariant:

$$\mathit{pathLength}(h, z, \mathbf{null}) = \mathit{pathLength}(h, t, \mathbf{null}).$$

Similar to the previous case, this invariant covers the scenario where the lists pointed by x and y are acyclic, as well as the one where they are cyclic. In the latter situation, the program does not terminate.

For the example in Figure 1c, the *divides* function is safe and a safety invariant is:

$$isPath(x, \mathbf{null}) \wedge isPath(z, \mathbf{null}) \wedge isPath(y, \mathbf{null}) \wedge isPath(y, z) \wedge isPath(x, w) \wedge \neg isNull(y) \wedge (pathLength(x, w) + pathLength(z, \mathbf{null}))\%pathLength(y, \mathbf{null}) = 0.$$

Additionally, the function terminates as witnessed by the ranking function $R(w) = pathLength(w, \mathbf{null})$.

Function $isCircular$ in Figure 1c is safe and terminating with the safety invariant: $pathLength(l, p) \wedge pathLength(p, q) \wedge isPath(q, p) \neq isPath(l, \mathbf{null})$, and lexicographic ranking function: $R(q, p) = (pathLength(q, \mathbf{null}), pathLength(q, p))$.

8 Experiments

To evaluate the applicability of our theory, we created a tool for verifying that heaps don't lie: Shakira [3]. We ran Shakira on a collection of programs manipulating singly linked lists. This collections includes the standard operations of traversal, reversal, sorting etc. as well as the motivational examples from Section 2. Each of the programs in this collection is annotated with correctness assertions and loop invariants, as well as the standard memory-safety checks. One of the programs (the motivational program from Figure 1b) used a non-linear loop invariant, but this did not require any special treatment by Shakira.

To generate VCs for each program, we generated a Hoare proof and then used CBMC 4.9 [4] to compute the strongest post-conditions for each Hoare triple using symbolic execution. The resulting VCs were solved using Glucose 4.0 [5]. As well as correctness and memory safety, these VCs proved that each loop annotation was genuinely a loop invariant. For four of the programs, we annotated loops with ranking functions and generated VCs to check that the loops terminated, thereby creating a total-correctness specification.

None of the proofs in our collection relied on assumptions about the shape of the heap beyond that it consisted of singly linked lists. In particular, our safety proofs show that the safe programs are safe even in the presence of arbitrary cycles and sharing between pointers.

We ran our experiments on a 4-core 3.30 GHz Core i5 with 8 GB of RAM. The results of these experiments are given in Table 1.

The top half of the table shows the aggregate results for the benchmarks in which the specifications held, i.e., the VCs were unsatisfiable. These "safe" benchmarks are divided into two categories: partial- and total-correctness proofs. Note that the total-correctness proofs involve solving more complex VCs – the partial correctness proofs solved 40 VCs in 5.9 s, while the total correctness proofs solved only 25 VCs in 9.6 s. This is due to the presence of ranking functions in the total-correctness proofs, which by necessity introduces a higher level of arithmetic complexity.

The bottom half of the table contains the results for benchmarks in which the VCs were satisfiable. Since the VCs were generated from a Hoare proof, their

	LOC	#VCs	Symex(s)	SAT(s)	C/E
Safe benchmarks (UNSAT VCs)					
SLL (safe)	236	40	18.2	5.9	—
SLL (termination)	113	25	14.7	9.6	—
Counterexamples (SAT VCs)					
CLL (nonterm)	38	14	6.9	1.6	3
Null-deref	165	31	13.6	3.0	3
Assertion Failure	73	11	3.5	0.7	3.5
Inadequate Invariant	37	4	4.9	1.2	6

Table 1: Experimental results

Legend:

LOC	Total lines of code
#VCs	Number of VCs
Symex(s)	Total time spent in symbolic execution to generate VCs
SAT(s)	Total time spent in SAT solver
C/E	Average counterexample size (number of nodes)

satisfiability only tells us that the purported proof is not in fact a real proof of the program's correctness. However, Shakira outputs models when the VCs are satisfiable and these can be examined to diagnose the cause of the proof's failure. For our benchmarks, the counterexamples fell into four categories:

- Non-termination due to cyclic lists.
- Null dereferences.
- A correctness assertion (not a memory-safety assertion) failing.
- The loop invariant being inadequate, either by being too weak to prove the required properties, or failing to be inductive.

A counterexample generated by Shakira is shown in Figure 7. This program is a variation on the motivational program from Figure 1c in which the programmer has tried to speed up the loop by unwinding it once. The result is that the program no longer terminates if the list contains a cycle whose size is exactly one, as shown in the counterexample found by Shakira.

These results show that discharging VCs written in SLH is practical with current technology. They further show that SLH is expressive enough to specify safety, termination and correctness properties for difficult programs. When the VCs require arithmetic to be done on list lengths, as is required when proving termination, the decision problem becomes noticeably more difficult. Our encoding is efficient enough that even when the VCs contain non-linear arithmetic on path lengths, they can be solved quickly by an off-the-shelf SAT solver.

9 Related Works

Research works on relating the shape of data structures to their numeric properties (e.g. length) follow several directions. For abstract interpretation based

```
int has_cycle(list l) {
  list p = l;
  list q = l→n;

  do {
    // Unwind loop to search
    // twice as fast!
    if (p != NULL) p = p→n;
    if (p != NULL) p = p→n;

    if (q != NULL) q = q→n;
    if (q != NULL) q = q→n;
    if (q != NULL) q = q→n;
    if (q != NULL) q = q→n;
  } while (p != q &&
           p != NULL &&
           q != NULL);

  return p == q;
}
```

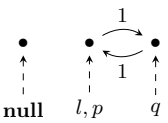

Counterexample heap leading to non-termination.

Fig. 7: A non-terminating program and the counterexample found by Shakira.

analyses, an abstract domain that captures both heap and size was proposed in [6]. The THOR tool [7, 8] implements a separation logic [9] based shape analysis and uses an off-the-shelf arithmetic analysis tool to add support for arithmetic reasoning. This approach is conceptually different from ours as it aims to separate the shape reasoning from the numeric reasoning by constructing a numeric program that explicitly tracks changes in data structure sizes. In [10], Boujjani et al. introduce the logic SLAD for reasoning about singly-linked lists and arrays with unbounded data, which allows to combine shape constraints, written in a fragment of separation logic, with data and size constraints. While SLAD is a powerful logic and has a decidable fragment, our main motivation for designing a new logic was its translation to SAT. A second motivation was the unrestricted sharing.

Other recent decidable logics for reasoning about linked lists were developed [2, 11–13, 10]. Piskac et al. provide a reduction of decidable separation logic fragments to a decidable first-order SMT theory [11]. A decision procedure for an alternation-free sub-fragment of first-order logic with transitive closure is described in [2]. Lahiri and Qadeer introduce the Logic of Interpreted Sets and Bounded Quantification (LISBQ) capable to express properties on the shape and data of composite data structures [14]. In [15], Brain et al. propose a decision procedure for reasoning about aliasing and reachability based on Abstract Conflict Driven Clause Learning (ACDCL) [16]. As they don't capture the lengths of lists, these logics are better suited for safety and less for termination proving.

10 Conclusions

We have presented the logic SLH for reasoning about potentially cyclic singly-linked lists. The main characteristics of SLH are the fact that it allows unrestricted sharing in the heap and can relate the structure of lists to their length, i.e. reachability constraints with numeric ones. As SLH is parametrised by the background arithmetic theory used to express the length of lists, we present its

instantiation SLH[$\mathcal{T}_{\mathcal{BV}}$] with the theory of bit-vector arithmetic and provide a way of efficiently deciding its validity via a reduction to SAT. We empirically show that SLH is both efficient and expressive enough for reasoning about safety and (especially) termination of list programs.

References

1. Immerman, N., Rabinovich, A.M., Reps, T.W., Sagiv, S., Yorsh, G.: The boundary between decidability and undecidability for transitive-closure logics. In: Computer Science Logic, 18th International Workshop, CSL 2004, 13th Annual Conference of the EACSL, Karpacz, Poland, September 20-24, 2004, Proceedings. (2004) 160–174
2. Itzhaky, S., Banerjee, A., Immerman, N., Nanevski, A., Sagiv, M.: Effectively-propositional reasoning about reachability in linked data structures. In: CAV. (2013) 756–772
3. Shakira: Hips Don't Lie (2006)
4. Clarke, E.M., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: TACAS. Springer (2004) 168–176
5. Audemard, G., Simon, L.: Predicting learnt clauses quality in modern SAT solvers. IJCAI'09, Morgan Kaufmann (2009) 399–404
6. Bouajjani, A., Dragoi, C., Enea, C., Sighireanu, M.: Abstract domains for automated reasoning about list-manipulating programs with infinite data. In: Verification, Model Checking, and Abstract Interpretation - 13th International Conference, VMCAI 2012, Philadelphia, PA, USA, January 22-24, 2012. Proceedings. (2012) 1–22
7. Magill, S., Tsai, M., Lee, P., Tsay, Y.: THOR: A tool for reasoning about shape and arithmetic. In: Computer Aided Verification, 20th International Conference, CAV 2008, Princeton, NJ, USA, July 7-14, 2008, Proceedings. (2008) 428–432
8. Magill, S., Tsai, M., Lee, P., Tsay, Y.: Automatic numeric abstractions for heap-manipulating programs. In: Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010. (2010) 211–222
9. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: LICS. (2002) 55–74
10. Bouajjani, A., Dragoi, C., Enea, C., Sighireanu, M.: Accurate invariant checking for programs manipulating lists and arrays with infinite data. In: ATVA. (2012)
11. Piskac, R., Wies, T., Zufferey, D.: Automating separation logic using SMT. In: CAV. (2013) 773–789
12. Yorsh, G., Rabinovich, A.M., Sagiv, M., Meyer, A., Bouajjani, A.: A logic of reachable patterns in linked data-structures. J.Log.Alg.Prog. **73**(1-2) (2007)
13. Madhusudan, P., Parlato, G., Qiu, X.: Decidable logics combining heap structures and data. In: POPL. (2011) 611–622
14. Lahiri, S.K., Qadeer, S.: Back to the future: revisiting precise program verification using SMT solvers. In: Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008. (2008) 171–182
15. Brain, M., David, C., Kroening, D., Schrammel, P.: Model and proof generation for heap-manipulating programs. In: ESOP. (2014) 432–452
16. D'Silva, V., Haller, L., Kroening, D.: Abstract conflict driven learning. In: POPL. (2013) 143–154